

Self-Organising Stochastic Encoders*

S P Luttrell

December 21, 2010

Abstract

The processing of mega-dimensional data, such as images, scales linearly with image size only if fixed size processing windows are used. It would be very useful to be able to automate the process of sizing and inter-connecting the processing windows. A stochastic encoder that is an extension of the standard Linde-Buzo-Gray vector quantiser, called a stochastic vector quantiser (SVQ), includes this required behaviour amongst its emergent properties, because it automatically splits the input space into statistically independent subspaces, which it then separately encodes.

Various optimal SVQs have been obtained, both analytically and numerically. Analytic solutions which demonstrate how the input space is split into independent subspaces may be obtained when an SVQ is used to encode data that lives on a 2-torus (e.g. the superposition of a pair of uncorrelated sinusoids). Many numerical solutions have also been obtained, using both SVQs and chains of linked SVQs: (1) images of multiple independent targets (encoders for single targets emerge), (2) images of multiple correlated targets (various types of encoder for single and multiple targets emerge), (3) superpositions of various waveforms (encoders for the separate waveforms emerge - this is a type of independent component analysis (ICA)), (4) maternal and foetal ECGs (another example of ICA), (5) images of textures (orientation maps and dominance stripes emerge).

Overall, SVQs exhibit a rich variety of self-organising behaviour, which effectively discovers the internal structure of the training data. This should have an immediate impact on “intelligent” computation, because it reduces the need for expert human intervention in the design of data processing algorithms.

1 Stochastic Vector Quantiser

1.1 Reference

Luttrell S P, 1997, *Mathematics of Neural Networks: Models, Algorithms and Applications*, Kluwer, Ellacott S W, Mason J C and Anderson I J (eds.), A theory of self-organising neural networks, 240-244.

*This talk was given at the Workshop on Self-Organising Systems - Future Prospects for Computing, 28-29 October 1999, Manchester, UK.

1.2 Objective Function

1.2.1 Mean Euclidean Distortion

$$D = \int d\mathbf{x} \Pr(\mathbf{x}) \sum_{\mathbf{y}} \Pr(\mathbf{y}|\mathbf{x}) \int d\mathbf{x}' \Pr(\mathbf{x}'|\mathbf{y}) \|\mathbf{x} - \mathbf{x}'\|^2 \quad (1)$$

- Encode then decode: $\mathbf{x} \longrightarrow \mathbf{y} \longrightarrow \mathbf{x}'$.
- \mathbf{x} = input vector; \mathbf{y} = code; \mathbf{x}' = reconstructed vector.
- Code vector $\mathbf{y} = (y_1, y_2, \dots, y_n)$ for $1 \leq y_i \leq M$.
- $\Pr(\mathbf{x})$ = input PDF; $\Pr(\mathbf{y}|\mathbf{x})$ = stochastic encoder; $\Pr(\mathbf{x}'|\mathbf{y})$ = stochastic decoder.
- $\|\mathbf{x} - \mathbf{x}'\|^2$ = Euclidean reconstruction error.

1.2.2 Simplify

$$\begin{aligned} D &= 2 \int d\mathbf{x} \Pr(\mathbf{x}) \sum_{\mathbf{y}} \Pr(\mathbf{y}|\mathbf{x}) \|\mathbf{x} - \mathbf{x}'(\mathbf{y})\|^2 \\ \mathbf{x}'(\mathbf{y}) &\equiv \int d\mathbf{x} \Pr(\mathbf{x}|\mathbf{y}) \mathbf{x} = \frac{\int d\mathbf{x} \Pr(\mathbf{x}) \Pr(\mathbf{y}|\mathbf{x}) \mathbf{x}}{\int d\mathbf{x} \Pr(\mathbf{x}) \Pr(\mathbf{y}|\mathbf{x})} \end{aligned} \quad (2)$$

- Do the $\int d\mathbf{x} \Pr(\mathbf{x}|\mathbf{y}) (\dots)$ integration.
- $\mathbf{x}'(\mathbf{y})$ = reconstruction vector.
- $\mathbf{x}'(\mathbf{y})$ is the solution of $\frac{\partial D}{\partial \mathbf{x}'(\mathbf{y})} = 0$, so it can be deduced by optimisation.

1.2.3 Constrain

$$\begin{aligned} \Pr(\mathbf{y}|\mathbf{x}) &= \Pr(y_1|\mathbf{x}) \Pr(y_2|\mathbf{x}) \dots \Pr(y_n|\mathbf{x}) \\ \mathbf{x}'(\mathbf{y}) &= \frac{1}{n} \sum_{i=1}^n \mathbf{x}'(y_i) \end{aligned} \quad (3)$$

- $\Pr(\mathbf{y}|\mathbf{x})$ implies the components (y_1, y_2, \dots, y_n) of \mathbf{y} are conditionally independent given \mathbf{x} .
- $\mathbf{x}'(\mathbf{y})$ implies the reconstruction is a superposition of contributions $\mathbf{x}'(y_i)$ for $i = 1, 2, \dots, n$.
- The stochastic encoder samples n times from the same $\Pr(y|\mathbf{x})$.

1.2.4 Upper Bound

$$\begin{aligned}
D &\leq D_1 + D_2 \\
D_1 &\equiv \frac{2}{n} \int d\mathbf{x} \Pr(\mathbf{x}) \sum_{y=1}^M \Pr(y|\mathbf{x}) \|\mathbf{x} - \mathbf{x}'(y)\|^2 \\
D_2 &\equiv \frac{2(n-1)}{n} \int d\mathbf{x} \Pr(\mathbf{x}) \left\| \mathbf{x} - \sum_{y=1}^M \Pr(y|\mathbf{x}) \mathbf{x}'(y) \right\|^2 \quad (4)
\end{aligned}$$

- D_1 is a stochastic vector quantiser with the vector code \mathbf{y} replaced by a scalar code y .
- D_2 is a non-linear (note $\Pr(y|\mathbf{x})$) encoder with a superposition term $\sum_{y=1}^M \Pr(y|\mathbf{x}) \mathbf{x}'(y)$.
- $n \rightarrow \infty$: the stochastic encoder measures $\Pr(y|\mathbf{x})$ accurately and D_2 dominates.
- $n \rightarrow 1$: the stochastic encoder samples $\Pr(y|\mathbf{x})$ poorly and D_1 dominates.

2 Analytic Optimisation

2.1 References

Luttrell S P, 1999, *Combining Artificial Neural Nets: Ensemble and Modular Multi-Net Systems*, 235-263, Springer-Verlag, Sharkey A J C (ed.), Self-organised modular neural networks for encoding data.

Luttrell S P, 1999, An Adaptive Network For Encoding Data Using Piecewise Linear Functions, *Proceedings of the 9th International Conference on Artificial Neural Networks (ICANN99)*, Edinburgh, 7-10 September 1999, 198-203.

2.2 Stationarity Conditions

2.2.1 Stationarity w.r.t. $\mathbf{x}'(y)$

$$n \int d\mathbf{x} \Pr(\mathbf{x}|y) \mathbf{x} = \mathbf{x}'(y) + (n-1) \int d\mathbf{x} \Pr(\mathbf{x}|y) \sum_{y'=1}^M \Pr(y'|\mathbf{x}) \mathbf{x}'(y') \quad (5)$$

- Stationarity condition is $\frac{\partial(D_1+D_2)}{\partial \mathbf{x}'(y)} = 0$.

2.2.2 Stationarity w.r.t. $\Pr(y' | \mathbf{x})$

$$\Pr(\mathbf{x}) \Pr(y | \mathbf{x}) \sum_{y'=1}^M (\Pr(y' | \mathbf{x}) - \delta_{y,y'}) \mathbf{x}'(y') \cdot \left(\frac{1}{2} \mathbf{x}'(y') - n \mathbf{x} + (n-1) \sum_{y''=1}^M \Pr(y'' | \mathbf{x}) \mathbf{x}'(y'') \right) = 0 \quad (6)$$

- Stationarity condition is $\frac{\partial(D_1+D_2)}{\partial \log \Pr(y|\mathbf{x})} = 0$, subject to $\sum_{y=1}^M \Pr(y | \mathbf{x}) = 1$.
- 3 types of solution: $\Pr(\mathbf{x}) = 0$ (trivial), $\Pr(y | \mathbf{x}) = 0$ (ensures $\Pr(y | \mathbf{x}) \geq 0$), and $\sum_{y'=1}^M (\Pr(y' | \mathbf{x}) - \delta_{y,y'}) (\dots) = 0$.

2.3 Circle

2.3.1 Input vector uniformly distributed on a circle

$$\begin{aligned} \mathbf{x} &= (\cos \theta, \sin \theta) \\ \int d\mathbf{x} \Pr(\mathbf{x}) (\dots) &= \frac{1}{2\pi} \int_0^{2\pi} d\theta (\dots) \end{aligned} \quad (7)$$

2.3.2 Stochastic encoder PDFs symmetrically arranged around the circle


$$\Pr(y | \theta) = p\left(\theta - \frac{2\pi y}{M}\right) \quad (8)$$

2.3.3 Reconstruction vectors symmetrically arranged around the circle

$$\mathbf{x}'(y) = r \left(\cos\left(\frac{2\pi y}{M}\right), \sin\left(\frac{2\pi y}{M}\right) \right) \quad (9)$$


2.3.4 Stochastic encoder PDFs overlap no more than 2 at a time

$$p(\theta) = \begin{cases} 1 & 0 \leq |\theta| \leq \frac{\pi}{M} - s \\ f(\theta) & \frac{\pi}{M} - s \leq |\theta| \leq \frac{\pi}{M} + s \\ 0 & |\theta| \geq \frac{\pi}{M} + s \end{cases} \quad (10)$$



2.3.5 Stochastic encoder PDFs overlap no more than 3 at a time

$$p(\theta) = \begin{cases} f_1(\theta) & 0 \leq |\theta| \leq -\frac{\pi}{M} + s \\ f_2(\theta) & -\frac{\pi}{M} + s \leq |\theta| \leq \frac{3\pi}{M} - s \\ f_3(\theta) & \frac{3\pi}{M} - s \leq |\theta| \leq \frac{\pi}{M} + s \\ 0 & |\theta| \geq \frac{\pi}{M} + s \end{cases} \quad (11)$$



2.4 2-Torus


2.4.1 Input vector uniformly distributed on a 2-torus

$$\begin{aligned} \mathbf{x} &= (\mathbf{x}_1, \mathbf{x}_2) \\ \mathbf{x}_1 &= (\cos \theta_1, \sin \theta_1) \\ \mathbf{x}_2 &= (\cos \theta_2, \sin \theta_2) \\ \int d\mathbf{x} \, \text{Pr}(\mathbf{x}) (\cdots) &= \frac{1}{4\pi^2} \int_0^{2\pi} d\theta_1 \int_0^{2\pi} d\theta_2 (\cdots) \end{aligned} \quad (12)$$

2.4.2 Joint encoding

$$\Pr(y|\mathbf{x}) = \Pr(y|\mathbf{x}_1, \mathbf{x}_2) \quad (13)$$


- $\Pr(y|\mathbf{x})$ depends jointly on \mathbf{x}_1 and \mathbf{x}_2 .
- Requires $n = 1$ to encode \mathbf{x} .
- For a given resolution the size of the codebook increases exponentially with input dimension.



2.4.3 Factorial encoding


$$\Pr(y|\mathbf{x}) = \begin{cases} \Pr(y|\mathbf{x}_1) & y \in Y_1 \\ \Pr(y|\mathbf{x}_2) & y \in Y_2 \end{cases} \quad (14)$$

- Y_1 and Y_2 are non-intersecting subsets of the allowed values of y .
- $\Pr(y|\mathbf{x})$ depends either on \mathbf{x}_1 or on \mathbf{x}_2 , but not on both at the same time.
- Requires $n \gg 1$ to encode \mathbf{x} .
- For a given resolution the size of the codebook increases linearly with input dimension.



2.4.4 Stability diagram

- Fixed n , increasing M : joint encoding is eventually favoured because the size of the codebook is eventually large enough.
- Fixed M , increase n : factorial encoding is eventually favoured because the number of samples is eventually large enough.
- Factorial encoding is encouraged by using a small codebook and sampling a large number of times.



3 Numerical Optimisation

3.1 References

Luttrell S P, 1997, to appear in *Proceedings of the Conference on Information Theory and the Brain*, Newquay, 20-21 September 1996, The emergence of dominance stripes and orientation maps in a network of firing neurons.

Luttrell S P, 1997, *Mathematics of Neural Networks: Models, Algorithms and Applications*, Kluwer, Ellacott S W, Mason J C and Anderson I J (eds.), A theory of self-organising neural networks, 240-244.

Luttrell S P, 1999, submitted to a special issue of *IEEE Trans. Information Theory on Information-Theoretic Imaging*, Stochastic vector quantisers.

3.2 Gradient Descent

3.2.1 Posterior probability with infinite range neighbourhood

$$\Pr(y|\mathbf{x}) = \frac{Q(\mathbf{x}|y)}{\sum_{y'=1}^M Q(\mathbf{x}|y')}$$

- $Q(\mathbf{x}|y) \geq 0$ is needed to ensure a valid $\Pr(y|\mathbf{x})$.
- This does not restrict $\Pr(y|\mathbf{x})$ in any way.

3.2.2 Posterior probability with finite range neighbourhood

$$\Pr(y|\mathbf{x}; y') \equiv \frac{Q(\mathbf{x}|y) \delta_{y \in \mathcal{N}(y')}}{\sum_{y'' \in \mathcal{N}(y')} Q(\mathbf{x}|y'')} \quad (15)$$

$$\Pr(y|\mathbf{x}) = \frac{1}{M} \sum_{y' \in \mathcal{N}^{-1}(y)} \Pr(y|\mathbf{x}; y') = \frac{1}{M} Q(\mathbf{x}|y) \sum_{y' \in \mathcal{N}^{-1}(y)} \frac{1}{\sum_{y'' \in \mathcal{N}(y')} Q(\mathbf{x}|y'')} \quad (16)$$

- $\mathcal{N}(y')$ is the set of neurons that lie in a predefined “neighbourhood” of y' .
- $\mathcal{N}^{-1}(y)$ is the “inverse neighbourhood” of y defined as $\mathcal{N}^{-1}(y) \equiv \{y' : y \in \mathcal{N}(y')\}$.
- Neighbourhood is used to introduce “lateral inhibition” between the firing neurons.
- This restricts $\Pr(y|\mathbf{x})$, but allows limited range lateral interactions to be used.

3.2.3 Probability leakage

$$\Pr(y|\mathbf{x}) \longrightarrow \sum_{y' \in \mathcal{L}^{-1}(y)} \Pr(y|y') \Pr(y'|\mathbf{x}) \quad (17)$$

- $\Pr(y|y')$ is the amount of probability that leaks from location y' to location y .
- $\mathcal{L}(y')$ is the “leakage neighbourhood” of y' .
- $\mathcal{L}^{-1}(y)$ is the “inverse leakage neighbourhood” of y defined as $\mathcal{L}^{-1}(y) \equiv \{y' : y \in \mathcal{L}(y')\}$.
- Leakage is to allow the network output to be “damaged” in a controlled way.
- When the network is optimised it automatically becomes robust with respect to such damage.

- Leakage leads to topographic ordering according to the defined neighbourhood.
- This restricts $\Pr(y|\mathbf{x})$, but allows topographic ordering to be obtained, and is faster to train.

3.2.4 Shorthand notation

$$\begin{aligned}
L_{y,y'} &\equiv \Pr(y|y') & P_{y,y'} &\equiv \Pr(y|\mathbf{x};y') \\
p_y &\equiv \sum_{y' \in \mathcal{N}^{-1}(y)} P_{y',y} & (L^T p)_y &\equiv \sum_{y' \in \mathcal{L}^{-1}(y)} L_{y',y} p_{y'} \\
\mathbf{d}_y &\equiv \mathbf{x} - \mathbf{x}'(y) & (L \mathbf{d})_y &\equiv \sum_{y' \in \mathcal{L}(y)} L_{y',y} \mathbf{d}_{y'} \\
(P L \mathbf{d})_y &\equiv \sum_{y' \in \mathcal{N}(y)} P_{y,y'} (L \mathbf{d})_{y'} & (P^T P L \mathbf{d})_y &\equiv \sum_{y' \in \mathcal{N}^{-1}(y)} P_{y',y} (P L \mathbf{d})_{y'} \\
e_y &\equiv \|\mathbf{x} - \mathbf{x}'(y)\|^2 & (L e)_y &\equiv \sum_{y' \in \mathcal{L}(y)} L_{y,y'} e_{y'} \\
(P L e)_y &\equiv \sum_{y' \in \mathcal{N}(y)} P_{y,y'} (L e)_{y'} & (P^T P L e)_y &\equiv \sum_{y' \in \mathcal{N}^{-1}(y)} P_{y',y} (P L e)_{y'} \\
\bar{\mathbf{d}} &\equiv \sum_{y=1}^M (L^T p)_y \mathbf{d}_y & \text{or } \bar{\mathbf{d}} &\equiv \sum_{y=1}^M (P L \mathbf{d})_y
\end{aligned} \tag{18}$$

- This shorthand notation simplifies the appearance of the gradients of D_1 and D_2 .
- For instance, $\Pr(y|\mathbf{x}) = \frac{1}{M} (L^T p)_y$.

3.2.5 Derivatives w.r.t. $\mathbf{x}'(y)$

$$\begin{aligned}
\frac{\partial D_1}{\mathbf{x}'(y)} &= -\frac{4}{n M} \int d\mathbf{x} \Pr(\mathbf{x}) (L^T p)_y \mathbf{d}_y \\
\frac{\partial D_2}{\mathbf{x}'(y)} &= -\frac{4(n-1)}{n M^2} \int d\mathbf{x} \Pr(\mathbf{x}) (L^T p)_y \bar{\mathbf{d}}
\end{aligned} \tag{19}$$

- The extra factor $\frac{1}{M}$ in $\frac{\partial D_2}{\mathbf{x}'(y)}$ arises because there is a $\sum_{y=1}^M (\dots)$ hidden inside the $\bar{\mathbf{d}}$.

3.2.6 Functional derivatives w.r.t. $\log Q(\mathbf{x}|y)$

$$\begin{aligned}
\delta D_1 &= \frac{2}{n M} \int d\mathbf{x} \Pr(\mathbf{x}) \sum_{y=1}^M \delta \log Q(\mathbf{x}|y) \left(p_y (L e)_y - (P^T P L e)_y \right) \\
\delta D_2 &= \frac{4(n-1)}{n M^2} \int d\mathbf{x} \Pr(\mathbf{x}) \sum_{y=1}^M \delta \log Q(\mathbf{x}|y) \left(p_y (L \mathbf{d})_y - (P^T P L \mathbf{d})_y \right) \cdot \bar{\mathbf{d}}
\end{aligned} \tag{20}$$

- Differentiate w.r.t. $\log Q(\mathbf{x}|y)$ because $Q(\mathbf{x}|y) \geq 0$.

3.2.7 Neural response model

$$Q(\mathbf{x}|y) = \frac{1}{1 + \exp(-\mathbf{w}(y) \cdot \mathbf{x} - b(y))} \quad (21)$$

- This is a standard “sigmoid” function.
- This restricts $\Pr(y|\mathbf{x})$, but it is easy to implement, and leads to results similar to the ideal analytic results.


3.2.8 Derivatives w.r.t. $\mathbf{w}(y)$ and $b(y)$

$$\frac{\partial D_1}{\partial \begin{pmatrix} b(y) \\ \mathbf{w}(y) \end{pmatrix}} = \frac{2}{nM} \int d\mathbf{x} \Pr(\mathbf{x}) \left(p_y(L\mathbf{e})_y - (P^T P L \mathbf{e})_y \right) (1 - Q(\mathbf{x}|y)) \begin{pmatrix} 1 \\ \mathbf{x} \end{pmatrix} \quad (22)$$

$$\frac{\partial D_2}{\partial \begin{pmatrix} b(y) \\ \mathbf{w}(y) \end{pmatrix}} = \frac{4(n-1)}{nM^2} \int d\mathbf{x} \Pr(\mathbf{x}) \left(p_y(L\mathbf{d})_y - (P^T P L \mathbf{d})_y \right) \cdot \bar{\mathbf{d}} (1 - Q(\mathbf{x}|y)) \begin{pmatrix} 1 \\ \mathbf{x} \end{pmatrix}$$


3.3 Circle

3.3.1 Training history



- $M = 4$ and $n = 10$ were used.
- The reference vectors $\mathbf{x}'(y)$ (for $y = 1, 2, 3, 4$) are initialised close to the origin.
- The training history leads to stationary $\mathbf{x}'(y)$ just outside the unit circle.


3.3.2 Posterior probabilities



- Each of the posterior probabilities $\Pr(y|\mathbf{x})$ (for $y = 1, 2, 3, 4$) is large mainly in a $\frac{\pi}{2}$ radian arc of the circle.
- There is some overlap between the $\Pr(y|\mathbf{x})$.

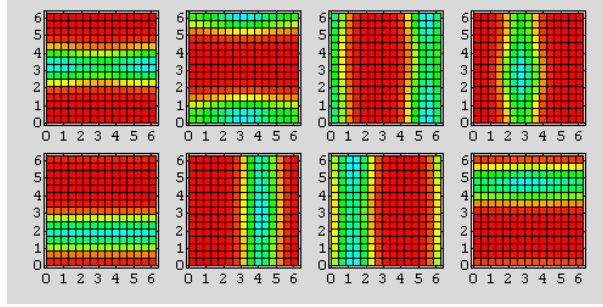
3.4 2-Torus

3.4.1 Posterior probabilities: joint encoding



- $M = 8$ and $n = 5$ were used, which lies inside the joint encoding region of the stability diagram.
- Each of the posterior probabilities $\Pr(y|\mathbf{x})$ is large mainly in a localised region of the torus.
- There is some overlap between the $\Pr(y|\mathbf{x})$.


3.4.2 Posterior probabilities: factorial encoding



- $M = 8$ and $n = 20$ were used, which lies inside the factorial encoding region of the stability diagram.
- Each of the posterior probabilities $\Pr(y|\mathbf{x})$ is large mainly in a collar-shaped region of the torus; half circle one way round the torus, and half the other way.
- There is some overlap between the $\Pr(y|\mathbf{x})$ that circle the same way round the torus.
- There is a localised region of overlap between a pair of $\Pr(y|\mathbf{x})$ that circle the opposite way round the torus.
- These localised overlap regions are the mechanism by which factorial encoding has a small reconstruction distortion.


3.5 Multiple Independent Targets

3.5.1 Training data



- The targets were unit height Gaussian bumps with $\sigma = 2$.
- The additive noise was uniformly distributed variables in $[0, 0.5]$.


3.5.2 Factorial encoding



- $M = 10$ and $n = 10$ were used.
- Each of the reference vectors $\mathbf{x}'(y)$ becomes large in a localised region.
- Each input vector causes a subset of the neurons to fire corresponding to locations of the targets.
- This is a factorial encoder because each neuron responds to only a subspace of the input.


3.6 Pair of Correlated Targets

3.6.1 Training data




- The targets were unit height Gaussian bumps with $\sigma = 1.5$.

3.6.2 Training history: joint encoding




- $M = 16$ and $n = 3$ were used.
- Each of the reference vectors $\mathbf{x}'(y)$ becomes large in a pair of localised regions.
- Each neuron responds to a small range of positions and separations of the pair of targets.
- The neurons respond *jointly* to the position and separation of the targets.

3.6.3 Training history: factorial encoding



- $M = \{16, 16\}$ and $n = \{20, 20\}$ were used; this is a 2-stage encoder.
- The second encoder uses as input the posterior probability output by the first encoder.
- The objective function is the sum of the separate encoder objective functions (with equal weighting given to each).
- The presence of the second encoder affects the optimisation of the first encoder via “self-supervision”.
- Each of the reference vectors $\mathbf{x}'(y)$ becomes large in a single localised region.


3.6.4 Training history: invariant encoding



- $M = \{16, 16\}$ and $n = \{3, 3\}$ were used; this is a 2-stage encoder.
- During training the ratio of the weighting assigned to the first and second encoders is increased from 1:5 to 1:40.
- Each of the reference vectors $\mathbf{x}'(y)$ becomes large in a single broad region.
- Each neuron responds only the position (and not the separation) of the pair of targets.
- The response of the neurons is *invariant* w.r.t. the separation of the targets.


3.7 Separating Different Waveforms

3.7.1 Training data



- This data is the superposition of a pair of waveforms plus noise.
- In each training vector the relative phase of the two waveforms is randomly selected.


3.7.2 Training history: factorial encoding



- $M = 10$ and $n = 20$ were used.
- Each of the reference vectors $\mathbf{x}'(y)$ becomes one or other of the two waveforms, and has a definite phase.
- Each neuron responds to only one of the waveforms, and then only when its phase is in a localised range.

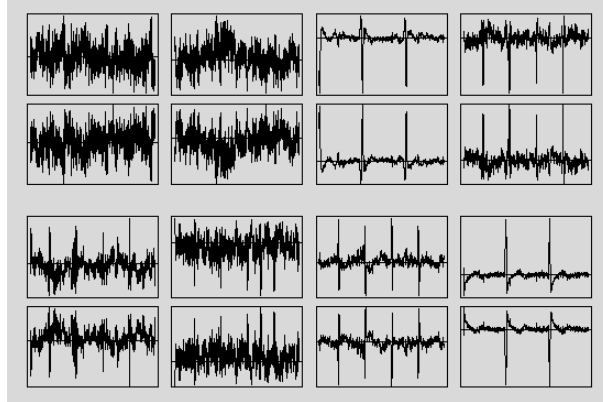
3.8 Maternal + Foetal ECG

3.8.1 Training data



- This data is an 8-channel ECG recording taken from a pregnant woman.
- The large spikes are the woman's heart beat.
- The noise masks the foetus' heartbeat.
- This data was whitened before training the neural network.

3.8.2 Factorial Encoding



- $M = 16$ and $n = 20$ were used.
- The results shown are $\mathbf{w}(y) \cdot \mathbf{x}$ computed for all neurons ($y = 1, 2, \dots, 8$) for each 8-dimensional input vector \mathbf{x} .
- After limited training some, but not all, of the neurons have converged.
- The broadly separated spikes indicate a neuron that responds to the mother's heartbeat.
- The closely separated spikes indicate a neuron that responds to the foetus' heartbeat.


3.9 Visual Cortex Network (VICON)

3.9.1 Training Data




- This is a Brodatz texture image, whose spatial correlation length is 5-10 pixels.

3.9.2 Orientation map



- $M = 30 \times 30$ and $n = 1$ were used.
- Input window size = 17×17 , neighbourhood size = 9×9 , leakage neighbourhood size = 3×3 were used.
- Leakage probability was sampled from a 2-dimensional Gaussian PDF, with $\sigma = 1$ in each direction.
- Each of the reference vectors $\mathbf{x}'(y)$ typically looks like a small patch of image.
- Leakage induces topographic ordering across the array of neurons
- This makes the array of reference vectors look like an “orientation map”.


3.9.3 Sparse coding



- The trained network is used to encode and decode a typical input image.

- Left image = input.
- Middle image = posterior probability. This shows “sparse coding” with a small number of “activity bubbles”.
- Right image = reconstruction. Apart from edge effects, this is a low resolution version of the input.

3.9.4 Dominance stripes



- Interdigitate a pair of training images, so that one occupies on the black squares, and the other the white squares, of a “chess board”.
- Preprocess this interdigitated image to locally normalise it using a finite range neighbourhood.
- $M = 100 \times 100$ and $n = 1$ were used.
- Input window size = 3×3 , neighbourhood size = 5×5 , leakage neighbourhood size = 3×3 were used.
- Leakage probability was sampled from a 2-dimensional Gaussian PDF, with $\sigma = 1$ in each direction.
- The dominance stripe map records for each neuron which of the 2 interdigitated images causes it to respond more strongly.
- The dominance stripes tend to run perpendicularly into the boundaries, because the neighbourhood window is truncated at the edge of the array.

4 References

Luttrell S P, 1997, to appear in *Proceedings of the Conference on Information Theory and the Brain*, Newquay, 20-21 September 1996, The emergence of dominance stripes and orientation maps in a network of firing neurons.

Luttrell S P, 1997, *Mathematics of Neural Networks: Models, Algorithms and Applications*, Kluwer, Ellacott S W, Mason J C and Anderson I J (eds.), A theory of self-organising neural networks, 240-244.

Luttrell S P, 1999, *Combining Artificial Neural Nets: Ensemble and Modular Multi-Net Systems*, 235-263, Springer-Verlag, Sharkey A J C (ed.), Self-organised modular neural networks for encoding data.

Luttrell S P, 1999, An Adaptive Network For Encoding Data Using Piecewise Linear Functions, *Proceedings of the 9th International Conference on Artificial Neural Networks (ICANN99)*, Edinburgh, 7-10 September 1999, 198-203.

Luttrell S P, 1999, submitted to a special issue of *IEEE Trans. Information Theory on Information-Theoretic Imaging*, Stochastic vector quantisers.